

Re: Design for historical data

Source:

<http://www.tech-archive.net/Archive/SQL-Server/microsoft.public.sqlserver.programming/2004-10/6481.html>

From: Zach Wells (*none_at_by.com*)

Date: 10/29/04

Date: Thu, 28 Oct 2004 21:06:18 -0500

Given the very brief overview, #2 is usually the way to go. If the three types of visits are really that, 3 separate types of visits, which store unique and different types of data, then they should be 3 separate tables. You can do all the formatting for the display on the front end when you need to make them all fit together for the UI.

As to "how many rows are too many"...well, that's really more based on your hardware. As an example, I'm currently running a database that has about 10 million rows of data in it that gets used daily (data only added/updated once a week though) and it runs fine on some crappy server that was just laying around.

More times than not, if there is a bottleneck while retrieving data, it is going to be poor table design and poor indexing. Hardware can be an issue (especially memory and disk access) but don't go spend 10k upgrading your server until you've optimized your queries.

Zach

ER Slansky wrote:

- > *Because my only real design experience is with Access, I don't know how*
- > *powerful SQL Server (2k) can be and the best way to design for it.*
- >
- > *Moving from Access because table of historical data is simply getting*
- > *too large. I will have to keep history on an individual potentially for*
- > *years (for doctors office) because the same doctor may not be the one to*
- > *see an individual each time they come in. All prior history must be*
- > *called up to review when making a diagnosis.*
- >
- > *There are also three types of visits with really nothing in common*
- > *except the patient, but the history of all three types of visits at*
- > *times must be displayed and other times only the history of a specific*
- > *type of visit displayed.*
- >
- > *So, and I apologize if this is a trivial question for experienced*
- > *designers, I see two options:*
- >

microsoft.public.sqlserver.programming: Re: Design for historical data

- > 1) *Combine everything in one table with indexes on visit type and patient – this would result in largest number of records in one table with the fields that are dissimilar containing nulls.*
- >
- > 2) *Three separate tables for each type of visit containing only the specific fields for that visit type and index off patient. To get a complete history, I would have to search all three tables for the patient, but when only needing history on one visit type, I would only have to go to that particular table.*
- >
- > *In Access, the design is currently one large table and to search, even with an index on patient, is taking longer as the number of records grows. I know SQL Server can handle it much better, but would it handle either option equally well? There are quite a few adds, and modifications to the "current" visit record for maybe 2 or 3 days after a visit, but really hardly ever any deletes.*
- >
- > *How many records is "too many" for SQL – or is there such a thing??*
- >
- > *Thanks,*
- > *ER Slansky*
- > *San Antonio, TX*
- >
- >
- > **** Sent via Developersdex <http://www.developersdex.com> ****
- > *Don't just participate in USENET...get rewarded for it!*